



Attorney Docket No.: 9209-7

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re: Lin Yu

Confirmation No. 9539

Serial No.: 09/896,179

Group Art Unit: 2171

Filed: June 29, 2001

Examiner: Te Y Chen

For: DEFERRED INDEX BUILDING SYSTEMS, METHODS AND COMPUTER
PROGRAM PRODUCTS FOR STORING TEMPORALLY SPACED APART
BURSTS OF DATA RECORDS IN A DATABASE

Date: August 16, 2004

RECEIVED

AUG 19 2004

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Technology Center 2100

APPELLANT'S BRIEF ON APPEAL UNDER 37 C.F.R. §1.192

Sir:

This Appeal Brief is filed pursuant to the "Notice of Appeal to the Board of Patent Appeals and Interferences" filed June 14, 2004.

Real Party In Interest

The real party in interest is assignee Trendium Inc., Sawgrass Lake Center, 13450 West Sunrise Blvd., Suite 200, Sunrise, Florida, 33323.

Related Appeals and Interferences

Appellant is aware of no appeals or interferences that would be affected by the present appeal.

Status of Claims

Appellant appeals the final rejection of Claims 1 - 42, which as of the filing date of this Brief remain under consideration. The claims at issue as finally rejected in the Office Action of March 24, 2004 are attached hereto as Appendix A.

08/18/2004 SZEWDIE1 00000038 09896179

01 FC:1402

330.00 OP

Status of Amendments

The following responses have been filed in the present case: An "Amendment" was filed December 19, 2003 in response to an Office Action mailed November 12, 2003. A "Request For Reconsideration After Final Action Pursuant To 37 CFR 1.116" was filed April 20, 2004 in response to a Final Office Action mailed March 24, 2004. An Advisory Action was mailed May 12, 2004. Claims 1 - 42 remain for consideration on the present appeal.

Summary of the Invention

Appellant's invention is directed to systems, methods and/or computer program products for storing temporally spaced apart bursts of data records in a database, by deferring building an index for a plurality of data records in a respective burst, until after storing the plurality of data records in the respective burst in the database. (Specification, page 5, lines 14 - 28). Thus, while the data burst is being received, little or no resources may need to be devoted to index building. Rather, index building may begin after termination of a data burst, when more resources may be available. In other embodiments of the invention, temporally spaced apart bursts of data records are received during a corresponding series of spaced apart time intervals. Deferred building of an index takes place by storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals, and beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

(Specification, page 6, lines 9 - 19 and FIG. 3).

Issue

Are claims 1 - 42 properly rejected under 35 U.S.C. §102(e) as being unpatentable over U. S. Patent No. 5,204,958 to Cheng et al. (hereinafter Cheng)?

Grouping of Claims

Claims 1 - 42 stand rejected under 35 U.S.C. §102(e) as being anticipated by Cheng. For purposes of this appeal, Claims 1, 2, 8, 14, 15, 21, 28, 29, 35, and 36 (Group I) may be considered as standing or falling together, and Claims 3 - 7, 9 - 13, 16 - 20, 22 - 27, 30 - 34, and 37 - 42 (Group II) may be considered as standing or falling together. Appellant respectfully submits that Claims 3 - 7, 9 - 13, 16 - 20, 22 - 27, 30 - 34, and 37 - 42 (Group II) are separately patentable from Claims 1, 2, 8, 14, 15, 21, 28, 29, 35, and 36 (Group I) because the Group II

claims include additional recitations directed towards storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals and beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

Argument

I. Introduction

Claims 1 - 42 stand rejected under 35 U.S.C. §102(e) as being anticipated by Cheng. Each of the claims of the present application stands rejected as anticipated under 35 U.S.C. § 102(e). Under 35 U.S.C. § 102, "a claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." M.P.E.P. § 2131 (quoting *Verdegaal Bros. v. Union Oil Co.*, 814 F.2d 628, 631, 2 U.S.P.Q.2d 1051, 1053 (Fed. Cir. 1987)). Stated another way, all material elements of a claim must be found in one prior art source. See *In re Marshall*, 198 U.S.P.Q. 344 (C.C.P.A. 1978). "Anticipation under 35 U.S.C. § 102 requires the disclosure in a single piece of prior art of each and every limitation of a claimed invention." *Apple Computer Inc. v. Articulate Sys. Inc.*, 57 U.S.P.Q.2d 1057, 1061 (Fed. Cir. 2000).

A finding of anticipation further requires that there must be no difference between the claimed invention and the disclosure of the cited reference as viewed by one of ordinary skill in the art. See *Scripps Clinic & Research Foundation v. Genentech Inc.*, 927 F.2d 1565, 1576, 18 U.S.P.Q.2d 1001, 1010 (Fed. Cir. 1991). In particular, the Court of Appeals for the Federal Circuit held that a finding of anticipation requires absolute identity for each and every element set forth in the claimed invention. See *Trintec Indus. Inc. v. Top-U.S.A. Corp.*, 63 U.S.P.Q.2d 1597 (Fed. Cir. 2002). Additionally, the cited prior art reference must be enabling, thereby placing the allegedly disclosed matter in the possession of the public. *In re Brown*, 329 F.2d 1006, 1011, 141 U.S.P.Q. 245, 249 (C.C.P.A. 1964). Thus, the prior art reference must adequately describe the claimed invention so that a person of ordinary skill in the art could make and use the invention.

Appellant respectfully submits that the pending claims are patentable over Cheng for at least the reason that this reference does not disclose all elements of the independent claims. The patentability of the pending claims is discussed in detail hereinafter.

II. The Group I Claims are Patentable over Cheng

Claims 1 - 8, 14 - 21, and 28 - 36 stand rejected under 35 U.S.C. §102(e) as being anticipated by Cheng. Independent Claims 1, 14, and 28 are method, system, and computer program product analogs of one another. For example, Claim 1 recites:

A method of storing temporally spaced apart bursts of data records in a database, comprising:

deferring building an index for a plurality of data records in a respective burst until after storing the plurality of data records in the respective burst in the database.

As noted in the present application, for example at page 2, line 34 - page 3, line 2:

Thus, while the data burst is being received, little or no resources may need to be devoted to index building. Rather, index building may begin after termination of a data burst, when more resources may be available.

Cheng relates to deferred index building, but in a different manner and for a different reason than described in the present application. In particular, Cheng relates to a "*System and Method for Efficiently Indexing and Storing a Large Database With High Data Insertion Frequency,*" as indicated in the title. Appellant respectfully submits that the Cheng disclosure relates to high frequency data insertion as explained at column 2, lines 40 - 50:

The present invention overcomes the above described disk I/O bottleneck associated with high frequency data insertion. It enables a computer system to perform high frequency inserts into the indexes of large databases with much less disk arm use, and therefore at much lower cost. In particular, the present invention defers index changes, and handles such updates to the stored log indexes in batches in a predefined order that matches the order in which indexes are stored on disk. As a result, the load imposed on disk devices is greatly reduced.

In sharp contrast, Claims 1, 14, and 28 relate to "bursty" data, such as was described in connection with Figure 1 of the present application, wherein "large amounts of data are received during a burst of time and no or relatively small amounts of data are received between the bursts of time" (Specification, page 2, lines 17 - 18).

Moreover, Cheng proposes a solution that is different from that which is recited in Claims 1, 14, and 28 as highlighted in Cheng's Summary of the Invention at column 2, line 60 - column 3, line 16:

SUMMARY OF THE INVENTION

In summary, the present invention is a database indexing methodology which allows low cost indexing of databases with very high insertion rates. A database index file is maintained by a computer system having primary random access memory and secondary memory. A record for each item added to the database is stored in a sequential file in secondary memory (disk storage) and an indexed pointer to the new

record is stored in a small B-tree stored in primary random access memory. The full index file for the database is a second, large B-tree stored in secondary memory. Leaf-nodes of the full index file are stored in packed, indexed order.

Periodically, a portion of the memory resident small B-tree is merged with a corresponding portion of the large B-tree by selecting a range of index values and retrieving from secondary memory all indexed pointers in the selected range of index values. The indexed pointers in the first B-tree in the selected range of index values are merged into the retrieved records, the resulting merged set of indexed pointers are stored in secondary memory in indexed order in a contiguous area of secondary memory at the tail of the large B-tree. As a result, the indexed pointers for newly added database records are written to secondary memory in batches, thereby accessing secondary memory very efficiently.

Accordingly, Cheng uses a system of secondary memory to allow high throughput data to be stored. Cheng, however, does not appear to describe or suggest bursts of data, deferring building an index for a plurality of records in a burst of data or deferring an index until after storing a plurality of data records in the respective burst in the database, as recited in Claims 1, 14, and 28.

In response to Appellant's argument set forth above, the Final Office Action mailed March 24, 2004 states at page 3, paragraph 4:

Applicant's arguments filed on 12/22/2003 have been fully considered but they are not persuasive.

Regarding applicant's arguments of 35 U.S.C. 102(3) rejection that: 1) "Cheng et al. uses a system of secondary memory to allow high throughput data to be stored"; 2) "Cheng et al. does not appear to describe or suggest burst of data, deferring building and [sic] index for a plurality of records in a burst of data, or deferring an index until after storing a plurality of data records in the respective burst in the database, as recited in Claim 1, 14 and 23." The examiner disagrees.

In reply to these arguments, the examiner points out that Cheng specifically disclosed a computer system to perform high frequency data insertion to resolve disk I/O bottleneck at a lower cost. [e.g., col. 2, lines 40-45] wherein, the high frequency data insertion is a burst of data storing processing. Furthermore, Cheng discloses his system handles I/O bottleneck high frequency data insertion by deferring index changes, and handles such updates to the stored log indexes in batches in a predefined order that matches the order in which indexes are stored on disk [e.g. col. 2, lines 45-50, the steps 300-302, Fig. 4A]. (Emphasis added).

Appellant respectfully submits that Cheng's disclosure does not relate to "bursts of data."

Bursts of data are defined, for example, in Claim 1:

A method of storing temporally spaced apart bursts of data records in a database, comprising:

deferring building an index for a plurality of data records in a respective burst until after storing the plurality of data records in the respective burst in the database.
(Emphasis added.)

Appellant notes that the passage cited by the Examiner (Cheng, column 2, lines 40-45) relates to "high frequency data insertion." The above-quoted passage from the Final Office Action mailed March 24, 2004 equates "high frequency data insertion" with "temporally spaced apart bursts of data records" as recited in Claim 1. In response, Appellant refers to the following passage from Cheng that provides examples of high frequency data insertion. Cheng states at column 1, line 53- column 2, line 30:

Consider as an example, a multi-user banking system running one hundred transactions per second. Each transaction updates a column value in a randomly chosen row from one of several different tables. Using a system with three primary tables, two of which are small enough to be maintained in primary memory and one of which is maintained in secondary memory, each transaction will require, on average, two I/O operations, for a total of 200 disk I/O operations per second. If each disk arm of a disk storage device can handle no more than 25 I/O operations per second, then eight disk arms would be required to handle this level of transactional activity. One can easily envision a use for indexed retrieval of log records by any of a number of keys: Account Identifier concatenated with a Timestamp, to answer questions by account holders about past transactions; Teller Identifier concatenated with a Timestamp, to make it easier to resolve cash drawer imbalances, etc. The duration of interest for such indexes can be quite long.

Now consider the added resources that would be needed to keep an Account-ID-Timestamp index on log records over a period of a month, using a standard B-tree indexed file. For those readers not familiar with B-tree indexes, note that these types of indexes are well known in the prior art, and will be explained in some detail below. For now, the only significance of the use of standard B-tree indexes is that the position of the index of the log record for each successive record is random--meaning that it can be at any position in the file. One hundred new log index entries are generated per second for the Account-ID-Timestamp index, eight hours per day, twenty working days per month. Thus there are about 57,600,000 new entries generated per month. In addition, each index entry will require at least ten bytes, resulting in an index table occupying about half a gigabyte of storage space. Clearly, most of the index table will not be memory resident. Since the position of each new inserted log record is random, it will typically require an average of one page read and one page write in order to insert the log record for each transaction. Thus, each index of this kind adds about 200 disk I/O operations per second, or an additional eight disk arms. (Emphasis added).

Thus, Cheng describes high frequency, continuous transactions. In the Advisory Action, the Examiner maintains that "storing temporally spaced apart bursts of data records in a data base" is recited in the preamble and therefore is insufficient to support patentability of the pending claims. Appellant respectfully disagrees. Claims 1, 14, and 28 refer to a respective burst of data records in the body of the claim, such as described in connection with Figure 1 of the present application, wherein "large amounts of data are received during a burst of time and no or relatively small amounts of data are received between the bursts of time"

(Specification, page 2, lines 17 - 18). "Bursty" is also defined in the Microsoft Computer Dictionary, 3rd Edition, as "transmitting data in spurts, or bursts, rather than in a continuous stream." In view of the above, it would not be obvious to use deferred index building for high frequency data in a system that uses bursty data.

Thus, while Appellant acknowledges that Cheng may provide solutions in systems and methods for efficiently indexing and storing a large database with a high data insertion frequency, Appellant submits that Cheng does not describe any solution for storing temporarily spaced apart bursts of data records in a database as recited in Claims 1, 14, and 28.

Appellant respectfully submits, therefore, that Cheng does not disclose each element of independent Claims 1, 14, and 28, and that Claims 2 - 8, 15 - 21, and 29 - 36 are patentable as depending from an allowable claim. Accordingly, Appellant respectfully requests that the rejection of Claims 1 - 8, 14 - 21, and 28 - 36, which includes all of the claims in Group I and Claims 3 - 7, 16 - 20, and 30 - 34 from Group II, be reversed.

III. The Group II Claims are Patentable over Cheng

Claims 9 - 13, 22 - 27, and 37 - 42 stand rejected under 35 U.S.C. §102(e) as being anticipated by Cheng. Independent Claims 9, 22, and 37 are method, system, and computer program product analogs of one another. For example, Claim 9 recites:

A method of storing, in a database, temporally spaced apart bursts of data records that are received during a corresponding series of spaced apart time intervals, the method comprising:

storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals; and

beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

As discussed above in connection with Claims 1, 14, and 28, Cheng does not appear to provide any description or suggestion of storing temporarily spaced apart bursts of data records that are received during a corresponding series of spaced apart time intervals. Nor does Cheng appear to describe or suggest storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals. Finally, Cheng does not appear to describe or suggest beginning to build the index for a corresponding one of the

spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals, as recited in Claims 9, 22, and 37.

With regard to dependent Claims 3 - 7, 16 - 20, and 30 - 34, these claims include all of the recitations from independent Claims 1, 14, and 28, respectively, and are, therefore, patentable over Cheng for at least the reasons stated above with respect to Claims 1, 14, and 28. Appellant further submits that dependent Claims 3 - 7, 16 - 20, and 30 - 34 are separately patentable for at least the reasons discussed above with respect to independent Claims 9, 22, and 37.

Appellant respectfully submits, therefore, that Cheng does not disclose each element of independent Claims 9, 22, and 37, and that Claims 10 - 13, 23 - 27, and 38 - 42 are patentable as depending from an allowable claim. Moreover, Appellant respectfully submits that dependent Claims 3 - 7, 16 - 20, and 30 - 34 are separately patentable for at least the reasons discussed above with respect to independent Claims 9, 22, and 37. Accordingly, Appellant respectfully requests that the rejection of Claims 3 - 7, 9 - 13, 16 - 20, 22 - 27, 30 - 34, and 37 - 42 (Group II) be reversed.

IV. Conclusion

In summary, Appellant respectfully submits that, with respect to Claims 1 - 42 (Groups I and II), Cheng does not disclose all elements of the independent Claims 1, 9, 14, 22, 28, and 37. Accordingly, Appellant respectfully requests reversal of the rejection of Claims 1 - 42 (Groups I and II) based on Cheng.

Respectfully submitted,



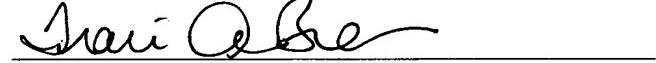
D. Scott Moore
Registration No. 42,011

Correspondence Address:

USPTO Customer No. 20792
Myers Bigel Sibley & Sajovec
Post Office Box 37428
Raleigh, North Carolina 27627
Telephone: 919/854-1400
Facsimile: 919/854-1401

CERTIFICATE OF MAILING

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Mail Stop Appeal Brief-Patents, Commissioner For Patents, Alexandria, VA 22313-1450, on August 16, 2004.


Traci A. Brown

APPENDIX A

1. A method of storing temporally spaced apart bursts of data records in a database, comprising:

deferring building an index for a plurality of data records in a respective burst until after storing the plurality of data records in the respective burst in the database.

2. A method according to Claim 1 wherein deferring building an index for a plurality of data records in a respective burst until after storing the plurality of data records in the respective burst in the database comprises:

deferring building an index for all the data records in a respective burst until after storing all the data records in the respective burst in the database.

3. A method according to Claim 2 wherein the temporally spaced apart bursts of data records are received during a corresponding series of spaced apart time intervals, and wherein deferring building an index for all the data records in a respective burst until after storing all the data records in the respective burst in the database further comprises:

storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals; and

beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

4. A method according to Claim 3 wherein beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

5. A method according to Claim 4 wherein building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals.

6. A method according to Claim 3 wherein storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals is performed by a first processor and beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals is performed by a second processor.

7. A method according to Claim 5 wherein storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals and building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals are performed alternately by a single processor.

8. A method according to Claim 1 wherein the database is an Indexed Sequential Access Method (ISAM) database.

9. A method of storing, in a database, temporally spaced apart bursts of data records that are received during a corresponding series of spaced apart time intervals, the method comprising:

storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals; and

beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

10. A method according to Claim 9 wherein beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

11. A method according to Claim 10 wherein building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals.

12. A method according to Claim 9 wherein storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals is performed by a first processor and beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals is performed by a second processor.

13. A method according to Claim 11 wherein storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals and building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals are performed alternately by a single processor.

14. A system for storing temporally spaced apart bursts of data records, comprising:

a database; and

means for deferring building an index for a plurality of data records in a respective burst until after storing the plurality of data records in the respective burst in the database.

15. A system according to Claim 14 wherein the means for deferring building an index for a plurality of data records in a respective burst until after storing the plurality of data records in the respective burst in the database comprises:

means for deferring building an index for all the data records in a respective burst until after storing all the data records in the respective burst in the database.

16. A system according to Claim 15 wherein the temporally spaced apart bursts of data records are received during a corresponding series of spaced apart time intervals, the means for deferring building an index for all the data records in a respective burst until after storing all the data records in the respective burst in the database further comprising:

means for storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals; and

means for beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

17. A system according to Claim 16 wherein the means for beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

means for building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

18. A system according to Claim 17 wherein the means for building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

means for building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals.

19. A system according to Claim 16 wherein the means for storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals is embodied in a first processor and the means for beginning is embodied in a second processor.

20. A system according to Claim 18 wherein the means for storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals and means for building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals are activated alternately in a single processor.

21. A system according to Claim 14 wherein the database is an Indexed Sequential Access Method (ISAM) database.

22. A system for storing, in a database, temporally spaced apart bursts of data records that are received during a corresponding series of spaced apart time intervals, the system comprising:

means for storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals; and

means for beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

23. A system according to Claim 22 wherein the means for beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

means for building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

24. A system according to Claim 23 wherein the means for building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

means for building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals.

25. A system according to Claim 22 wherein the means for storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals is embodied in a first processor and the means for beginning to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals is embodied in a second processor.

26. A system according to Claim 24 wherein the means for storing the spaced apart bursts of data records in the database during the corresponding series of spaced apart

time intervals and means for building the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals are activated alternately in a single processor.

27. A system according to Claim 22 further comprising the database.

28. A computer program product for storing temporally spaced apart bursts of data records in a database, the computer program product comprising a computer-readable storage medium having computer-readable program code embodied in the medium, the computer-readable program code comprising:

computer-readable program code that is configured to defer building an index for a plurality of data records in a respective burst until after storing the plurality of data records in the respective burst in the database.

29. A computer program product according to Claim 28 wherein the computer-readable program code that is configured to defer building an index for a plurality of data records in a respective burst until after storing the plurality of data records in the respective burst in the database comprises:

computer-readable program code that is configured to defer building an index for all the data records in a respective burst until after storing all the data records in the respective burst in the database.

30. A computer program product according to Claim 28 wherein the temporally spaced apart bursts of data records are received during a corresponding series of spaced apart time intervals, the computer-readable program code that is configured to defer building an index for all the data records in a respective burst until after storing all the data records in the respective burst in the database further comprising:

computer-readable program code that is configured to store the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals; and

computer-readable program code that is configured to begin to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

31. A computer program product according to Claim 30 wherein the computer-readable program code that is configured to begin to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

computer-readable program code that is configured to build the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

32. A computer program product according to Claim 31 wherein the computer-readable program code that is configured to build the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

computer-readable program code that is configured to build the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals.

33. A computer program product according to Claim 30 wherein the computer-readable program code that is configured to store the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals is configured to execute on a first processor and the computer-readable program code that is configured to begin to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals is configured to execute on a second processor.

34. A computer program product according to Claim 32 wherein the computer-readable program code that is configured to store the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals and the computer-readable program code that is configured to build the index for the corresponding one of the

spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals are configured to execute alternately on a single processor.

35. A computer program product according to Claim 28 wherein the database is an Indexed Sequential Access Method (ISAM) database.

36. A computer program product according to Claim 28 further comprising:
computer-readable program code that is configured to provide the database.

37. A computer program product for storing, in a database, temporally spaced apart bursts of data records that are received during a corresponding series of spaced apart time intervals, the computer program product comprising a computer-readable storage medium having computer-readable program code embodied in the medium, the computer-readable program code comprising:

computer-readable program code that is configured to store the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals;
and

computer-readable program code that is configured to begin to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

38. A computer program product according to Claim 37 wherein the computer-readable program code that is configured to begin to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

computer-readable program code that is configured to build the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals.

39. A computer program product according to Claim 38 wherein the computer-readable program code that is configured to build the index for the corresponding one of the

spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals comprises:

computer-readable program code that is configured to build the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals.

40. A computer program product according to Claim 37 wherein the computer-readable program code that is configured to store the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals is configured to execute on a first processor and the computer-readable program code that is configured to begin to build the index for a corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals is configured to execute on a second processor.

41. A computer program product according to Claim 39 wherein the computer-readable program code that is configured to store the spaced apart bursts of data records in the database during the corresponding series of spaced apart time intervals and the computer-readable program code that is configured to build the index for the corresponding one of the spaced apart bursts after expiration of the corresponding one of the series of spaced apart time intervals and prior to beginning a next one of the series of spaced apart time intervals are configured to execute alternatingly on a single processor.

42. A computer program product according to Claim 37 further comprising computer-readable program code that is configured to provide the database.